

APPLICATION FOR UNITED STATES LETTERS PATENT
FOR
**HYBRID OPTICAL BURST SWITCHING WITH FIXED TIME SLOT
ARCHITECTURE**

INVENTOR(S): **MANAV MISHRA**
CHRISTIAN MACIOCCO

PREPARED BY:

BLAKELY SOKOLOFF TAYLOR & ZAFMAN LLP
12400 WILSHIRE BOULEVARD, 7TH FLOOR
LOS ANGELES, CALIFORNIA 90025
(206) 292-8600

Express Mail Certificate of Mailing

"Express Mail" mailing label number: EV320119563US
Date of Deposit: December 22, 2003

*I hereby certify that this paper or fee is being deposited with the
United States Postal Service "Express Mail Post Office to Addressee"
service under 37 CFR 1.10 on the date indicated above and is
addressed to the Mail Stop Patent Application, Commissioner for
Patents, P.O. Box 1450, Alexandria, VA 22313-1450.*

Adrian Villarreal

(Typed or printed name of person mailing paper or fee)



(Signature of person mailing paper or fee)

December 22, 2003
(Date signed)

HYBRID OPTICAL BURST SWITCHING WITH FIXED TIME SLOT

ARCHITECTURE

CROSS REFERENCE TO RELATED APPLICATIONS

[0001] The present application is related to U.S. Patent Application Serial No.
5 10/674,650, filed 09/30/03, entitled "OPTICAL-SWITCHED (OS) NETWORK TO OS
NETWORK ROUTING USING EXTENDED BORDER GATEWAY PROTOCOL,"
and U.S. Patent Application Serial No. 10/636,062, filed 08/06/03, entitled
"RESERVATION PROTOCOL SIGNALING EXTENSIONS FOR OPTICAL
SWITCHED NETWORKS," both of which are assigned to the Assignee of the present
10 application.

TECHNICAL FIELD

[0002] This disclosure relates generally to optical networks, and in particular
but not exclusively, relates to optical burst switching networks.
15

BACKGROUND INFORMATION

[0003] Transmission bandwidth demands in telecommunication networks (e.g.,
the Internet) appear to be ever increasing and solutions are being sought to support this
bandwidth demand. One solution to this problem is to use fiber-optic networks, where
20 wavelength-division-multiplexing ("WDM") technology is used to support the ever-
growing demand in optical networks for higher data rates.

[0004] Conventional optical switching networks typically use wavelength
routing techniques, which require optical-to-electrical-to-optical ("O-E-O") conversions

of optical signals at each optical switching node. O-E-O conversions at each switching node in the optical network is not only a very slow operation (typically about ten milliseconds), but it is very costly, and potentially creates a traffic bottleneck for the optical switched network. In addition, the current optical switching technologies cannot efficiently support “bursty” traffic that is often experienced in packet based communication applications (e.g., the Internet).

[0005] Recently, optical burst switching (“OBS”) schemes have emerged as a promising solution to support high-speed bursty data traffic over WDM optical networks. The OBS scheme offers a practical opportunity between the current circuit-switching and the emerging all optical packet switching technologies. It has been shown that under certain conditions, the OBS scheme achieves high-bandwidth utilization and class-of-service (“CoS”) by elimination of electronic bottlenecks as a result of O-E-O conversions occurring at switching nodes, and by using a one-way end-to-end bandwidth reservation scheme with variable time slot duration provisioning scheduled by the ingress nodes.

[0006] However, current signaling protocols, such as Just-In-Time (“JIT”) signaling, Just-Enough-Time (“JET”) signaling, or other one-way reservation protocols, employed to reserve network and device resources and to setup and teardown burst paths over OBS networks are complex and resource intensive. As such, expensive hardware, including powerful network processors and vast amounts of memory for tracking reservation protocol states, are required in each switching node of an OBS network. Despite the complexity of these signaling protocols, low loss of optical packets cannot be guaranteed. Furthermore, JIT and JET signaling are intimately wrapped up with the

task of scheduling data packets for transmission. As such, JIT and JET are inflexible solutions that prevent easy modification of the signaling protocol or scheduling algorithm used at a given time on an OBS network.

BRIEF DESCRIPTION OF THE DRAWINGS

[0007] Non-limiting and non-exhaustive embodiments of the present invention are described with reference to the following figures, wherein like reference numerals refer to like parts throughout the various views unless otherwise specified.

5 [0008] FIG. 1 is a block diagram illustrating a hybrid optical burst switching network for switching fixed length optical cells, in accordance with an embodiment of the present invention.

[0009] FIG. 2 is a block diagram illustrating functional blocks of an ingress node for generating optical bursts having fixed length optical cells, in accordance with
10 an embodiment of the present invention.

[0010] FIG. 3 is a flow chart illustrating a process for generating and transmitting fixed length optical cells over a hybrid optical burst switching network, in accordance with an embodiment of the present invention.

[0011] FIG. 4 is a diagram illustrating how to slice data streams into data slices
15 and to insert the data slices into fixed length optical cells for transmission as an optical burst, in accordance with an embodiment of the present invention.

[0012] FIG. 5 is a block diagram illustrating hardware blocks of an edge node for generating optical bursts having fixed length optical cells, in accordance with an embodiment of the present invention.

DETAILED DESCRIPTION

[0013] Embodiments of a system and method for a hybrid optical burst switching (“OBS”) network using fixed length optical cells are described herein. In the following description numerous specific details are set forth to provide a thorough understanding of embodiments of the invention. One skilled in the relevant art will recognize, however, that the invention can be practiced without one or more of the specific details, or with other methods, components, materials, etc. In other instances, well-known structures, materials, or operations are not shown or described in detail to avoid obscuring aspects of the invention.

[0014] Reference throughout this specification to “one embodiment” or “an embodiment” means that a particular feature, structure, or characteristic described in connection with the embodiment is included in at least one embodiment of the present invention. Thus, the appearances of the phrases “in one embodiment” or “in an embodiment” in various places throughout this specification are not necessarily all referring to the same embodiment. Furthermore, the particular features, structures, or characteristics may be combined in any suitable manner in one or more embodiments.

[0015] In short, embodiments of the present invention combine the high bandwidth capability of an OBS network with the simplicity of cell switching techniques to provide a hybrid OBS network. Instead of switching on the boundaries of variable length optical bursts, techniques are described herein for switching on the boundaries of fixed length optical cells without using a dynamic reservation algorithm. Rather, embodiments of the present invention replace the dynamic reservation algorithm with time slots (a.k.a. fixed length optical cells) assigned to edge nodes. Embodiments of the

present invention take advantage of less hardware intensive signaling protocols, such as Resource Reservation Protocol—Traffic Engineering (“RSVP-TE”) or the like, to establish coarse-grain (e.g., minutes, hours, days, or longer) optical paths across the hybrid OBS network. Furthermore, embodiments of the present invention describe an
5 ingress node architecture that isolates the signaling protocols for establishing optical paths from scheduling algorithms that determine the order in which inbound data streams are serviced. Embodiments of the present invention tradeoff overall data throughput and scalability for the simplicity derived from fixed length optical cells, assigned time slots, and other advantages described herein. The techniques described
10 herein are well suited for small optical networks. These and other embodiments are described in detail below.

[0016] FIG. 1 is a block diagram illustrating a hybrid optical burst switching (“OBS”) network 100 for switching fixed length optical cells, in accordance with an embodiment of the present invention. FIG. 1 further illustrates local area networks
15 (“LANs”) 105 and a wide area network (“WAN”) 107 communicatively coupled to hybrid OBS network 100 to send and to receive data streams thereto/therefrom.

[0017] Hybrid OBS network 100 is a type of optical switching network. The illustrated embodiment of hybrid OBS network 100 includes ingress nodes 110, egress nodes 115, and switching nodes 120. Hybrid OBS network 100 can include more or less
20 ingress nodes 110, egress nodes 115, or switching nodes 120 that are interconnected via switching nodes 120, as illustrated. Ingress nodes 110 and egress nodes 115 are also referred to as edge nodes in that they logically reside at the edge of hybrid OBS network

100. The edge nodes provide an interface between hybrid OBS network 100 and LANs 105 and WAN 107 and switching nodes 120.

[0018] LANs 105 and WAN 107 can be either optical networks or electrical networks. Therefore, in some embodiments, ingress nodes 110 may perform optical-to-electrical (“O-E”) conversions of inbound data streams. Ingress nodes 110 may further include electronic memory to buffer the inbound data streams as data blocks until the data blocks are ready to be transmitted to one of switching nodes 120 for transmission to one of egress nodes 115. In some embodiment, ingress nodes 110 may perform electrical-to-optical (“O-E”) conversion of the buffered data blocks prior to transmitting the data block as an optical burst to one of switching nodes 120.

[0019] Egress nodes 115 are implemented with optical switching modules that are configured to receive optical bursts from switching nodes 120 and route the data streams contained therein to LANs 105, WAN 107, or other external networks. In one embodiment, egress nodes 115 perform O-E-O conversions of received optical bursts. To perform the O-E-O conversions, egress nodes 115 may include electronic memory to buffer the data bursts as data blocks prior to forwarding the data blocks as data streams to an external network. Although ingress nodes 110 and egress nodes 110 have been illustrated as separate entities, it should be appreciated that ingress nodes 110 and egress nodes 115 illustrate separate functionalities that may be implemented within a single hardware device or edge node.

[0020] Switching nodes 120 are implemented with optical switching modules that are each configured to receive an optical burst from either one of ingress nodes 110 or another switching node 120 and appropriately route the optical burst to another

switching node 120 or to one of egress nodes 115. Switching nodes 120 switch optical bursts on the boundaries of fixed length optical cells per carrier wavelength, where each carrier wavelength can have its own fixed optical cell length thus providing traffic differentiation capabilities, as opposed to variable length optical bursts. Switching on
5 fixed length time slots (e.g., optical cells) simplifies the task of reserving resources along an optical path within hybrid OBS network 100. Thus, less expensive hardware and simpler routing/signaling protocols may be used.

[0021] Switching nodes 120 determine where to route an optical burst via control signals issued over a control plane, which establish optical paths through hybrid
10 OBS network 100 (discussed in detail below). In one embodiment, the optical paths are established in advance of ingress nodes 110 receiving data streams from the external networks. In one embodiment, these optical paths are established during a setup and configuration period of hybrid OBS network 110 and persist until a management station (not illustrated) coupled to the control plane reconfigures the optical paths. Furthermore,
15 in some embodiments, the management station assigns fixed slot sizes (e.g., optical cell sizes) to each carrier wavelength on hybrid OBS network 100 during this initial setup and configuration period. The fixed slot sizes assigned to each carrier wavelength may be used by the edge nodes (e.g., ingress nodes 110 and egress nodes 115) to appropriately assign traffic types for transmission on each of the carrier wavelengths.
20 The management station also assigns positions to each optical cell relative to a time offset for specific edge nodes (i.e. an edge node can only transmit during its assigned time slots or optical cells).

[0022] Hybrid OBS network 100 may use wavelength-division-multiplexing (“WDM”) to multiplex inbound data streams onto multiple wavelengths. In some embodiments, the control signals are optical control signals propagated on pre-selected carrier wavelengths out-of-band (“OOB”) of the data plane propagating optical bursts

5 carrying the data streams. In some embodiments, the control plane is a separate electrical network, such as an Ethernet. The control plane is often referred to as the slow path, since it does not support the high bandwidths available in the data plane, which is often referred to as the fast path. In yet other embodiments, the control signals are propagated in band on the data plane using different encoding schemes, such as different

10 modulation formats, using a recognizable header, or the like. It should be appreciated that embodiments of the present invention should not be limited by the techniques used to convey control signals for establishing optical paths.

[0023] FIG. 2 is a block diagram illustrating functional blocks of an ingress node 200 for generating optical bursts having fixed length optical cells, in accordance

15 with an embodiment of the present invention. Ingress node 200 is an exemplary embodiment of one of ingress nodes 110. The illustrated embodiment of ingress node 200 includes a buffer 205, a stream slicer 210, a header pre-append block 215, a scheduler 220, and a burst transmit block 225.

[0024] In short, ingress node 200 operates as follows. Ingress node 200

20 receives data streams 230 and temporarily stores data streams 230 as data blocks 235. Stream slicer 210 slices data blocks 235 into data slices 240. Header pre-append block 215 generates slice headers 245 for each data slice 240 and appends slice headers 245 thereto. Scheduler 220 schedules each data slice 240 with its pre-appended slice header

245 to be transmitted during fixed time slots, illustrated as optical cells 260A and 260B (collectively referred to as optical cells 260). Scheduler 220 may schedule data slices 240 for one or more data blocks 235 at once to generate one or more complete optical bursts, such as optical bursts 270A and 270B (collectively referred to as optical bursts 270). Once an entire optical burst has been prepared and scheduled, burst transmit block 225 converts the scheduled data slices 240 from the electrical realm to the optical realm and launches the optical bursts 270A and 270B onto a fiber link 250 for transmission to one of switching nodes 120. As illustrated in FIG. 2, each of optical bursts 270A and 270B may be modulated on different carrier wavelength. Furthermore, each carrier wavelength may transmit optical cells having a different length, though the length of each optical cell for a given carrier wavelength is constant throughout hybrid OBS network 100. It is important to note that each optical burst 270 may vary in overall length, but only increments equal to its optical cell size. Thus, optical bursts 270 have boundaries aligned with fixed length time slots (i.e., optical cells 260).

[0025] FIG. 3 is a flow chart illustrating a process 300 for generating and for transmitting data slices 240 within optical cells 260, in accordance with an embodiment of the present invention.

[0026] In a process block 305, the management station (not illustrated) executes a signaling protocol to establish optical paths through hybrid OBS network 100. In one embodiment, a signaling protocol called Resource Reservation Protocol-Traffic Engineering ("RSVP-TE"), with extensions appropriate for hybrid OBS network 100, is executed to reserve resources and establish optical paths. RSVP-TE may be used in connection with MultiProtocol Label Switching ("MPLS") to route control information

(e.g., routing labels) to switching nodes 120 to support Generalized MultiProtocol Label Switching ("GMPLS"). More information on GMPLS signaling with RSVP-TE extensions can be found at <http://www.ietf.org/rfc.html> (RFC 3473). RSVP-TE is a relatively simple signaling protocol compared to Just-In-Time ("JIT") and Just-Enough-Time ("JET"). As such, RSVP-TE can be implemented with less expensive control plane hardware in each of switching nodes 120. It should be appreciated that any number of signaling protocols for establishing optical paths across hybrid OBS network 100 and reserving resources may be implemented in connection with embodiments of the present invention. Thus, the architecture of ingress node 200 described herein is advantageously flexible in terms of coarse grain signaling protocols implemented in connection therewith.

[0027] In the case of JIT and JET, scheduling data streams for transmission across an optical network is complex and difficult task that is susceptible to wasteful retransmissions due to scheduling conflicts. When an inbound data stream begins to pour into an ingress node from an external network, JIT and JET immediately commence establishing an optical path and reserving resources along the path for an estimated period of time. However, if the inbound data stream does not completely arrive until later than expected, transmission of an optical burst containing the data stream will not complete transmission across the optical network within the estimated period of time, for which resources have been reserved. Resultantly, the subsequent optical burst will be delayed causing failed transmissions and required retransmissions. Known OBS networks simply cannot guarantee low loss rates. However, embodiments of the present invention do not suffer from this Achilles' heal. Rather, the management

station assigns optical cell sizes per carrier wavelength, as well as, the position of each optical cell to the edge nodes. These assignments prevent a node from transmitting an optical cell during time slot not assigned to the node. Since each optical cell 260 of an optical bursts 270 are transmitted during fixed time slots and only during these time
5 slots, data collisions do not occur. Therefore, embodiments of the present invention can guarantee low loss rates.

[0028] Optical paths may be established during an initial setup and configuration of hybrid OBS network 100. These optical paths are coarse grain optical paths that may persist for minutes, hours, days, or as long as is desired, as opposed to
10 sub-microsecond optical paths. In some embodiment, new optical paths may be established anytime and currently established paths liquidated under the control of the management station. However, embodiments of the present invention need not setup and tear down a new optical path for each optical burst 270 launched into hybrid OBS network 100. Using optical paths that persist over many optical bursts 270 reduces
15 latency that occurs at ingress nodes 110, as optical bursts 270 need not wait for an optical path to be established. Thus, embodiments of the present invention may be used in connection with “light weight” signaling protocols for establishing optical paths and reserving resources. This coarse grain optical path switching further provide simplicity of implementation at the expense of dynamic path switching and efficiency.

20 [0029] In a process block 310, ingress node 200 receives one or more data streams 230. As data streams 230 pour into ingress node 200 from external networks (e.g., LANs 105 or WAN 107), the data streams are temporarily stored within buffer 205 as data blocks 235. In one embodiment, ingress node 200 buffers an entire data stream

230 before commencing transmission of the data stream 230 as an optical burst 270 over hybrid OBS network 100. Thus, if ingress node 200 determines that a complete data block 235 has not yet been received, ingress node 200 continues to wait in a process block 320. In some embodiments, ingress node 200 may wait for several data streams 230 to be completely buffered in buffer 205 prior to generating one or more optical bursts 270.

[0030] In a process block 325, stream slicer 210 slices data blocks 235 into data slices 240. The lengths of the data slices 240 may vary based on hardware constraints, network optimizations, and the like. Data slices 240 may all have the same length with the possible exception of one data slice 240 per optical burst 270 to accommodate a burst header (discussed more below) or stream slicer 210 may scale the size of data slices 240 according to the length of optical cells 260. For example, a data slice 240A may be sliced by stream slicer 210 such that data slice 240A spans the entire length of a single optical cell 260A with only a single slice header 245A.

[0031] In a process block 330, header pre-append block 215, generates a slice header 245 for each data slice 240 chopped by stream slicer 210. Slice headers 245 are used by down stream entities, such as egress node 115 or a destination user, to reassemble data slices 240 into data blocks 235 or data streams 230. For example, if data slices corresponding to data stream 230A (and therefore data block 235A) are scheduled into different optical bursts 270, which in turn take different optical paths to a common destination, the data slices of data stream 230A may arrive at the destination out of order. Slice headers 245 contain ordering information to enable the destination to

reassemble the data slices into data stream 230A without requiring retransmission thereof.

5 **[0032]** In one embodiment, slice headers 245 each include a fragment identifier (“ID”) and a data stream ID. The fragment ID numbers the data slices of a particular data stream 230 so that the data stream 230 can be reassembled, as discussed above. The data stream ID identifies to which data stream 230 the data slice belongs. Including the data stream ID frees scheduler 220 to multiplex or schedule multiple data streams 230 onto one or more optical bursts 270 in any order desired to optimize bandwidth utilization.

10 **[0033]** In a process block 335, scheduler 220 schedules data slices 240 into optical cells 260 of one or more optical bursts 270. It should be appreciated that scheduler 220 can schedule data slices from one data stream 230 at a time or schedule data slices from multiple data streams 230 at a time. Furthermore, scheduler 220 can schedule data slices from a single data stream 230 into a single optical burst 270 or into
15 multiple optical bursts 270. In the case where multiple optical bursts 270 are used, optical bursts 270 may be launched concurrently on separate carrier wavelengths (e.g., optical bursts 270A and 270B), launched in a staggered or time delay fashion, or launched one after another on the same carrier wavelength.

20 **[0034]** Embodiments of the present invention separate the task of scheduling slices into optical cells 260 from the tasks of establishing optical paths and reserving resources. As such, any desirable scheduling algorithm may be implemented, such as round robin scheduling, deficit round robin (“DRR”) scheduling, weighted round robin (“WRR”) scheduling, or the like. Implementation of priority based scheduling

algorithms, such as WRR, enables guaranteed quality of service (“QoS”). The particular scheduling algorithm selected may be based on anticipated traffic profiles and demands. Furthermore, the scheduling algorithm may be adjusted or changed on the fly (in real-time) to respond to variable or bursty traffic patterns. Thus, it should be appreciated that
5 separating the tasks of scheduling from routing and resource reservation results in increased flexibility and responsiveness to traffic demands.

[0035] Embodiments of the present invention transmit data streams 230 over hybrid OBS network 100 within optical bursts 270 formed of a variable number of fixed length optical cells 260. As such, hybrid OBS network 100 executes an optical cell-
10 switching scheme. As is illustrated, data slices 240 are transmitted onto hybrid OBS network 100 within optical cells 260. Each optical cell 260 represents a fixed length time slot. Optical cells 260 have a uniform or constant length through hybrid OBS network 100 for each carrier wavelength (e.g., λ_1 or λ_2), but may vary from one carrier wavelength to the next. For example, FIG. 2 illustrates optical cells 260A of carrier
15 wavelength λ_1 as being twice as long as optical cells 260B of carrier wavelength λ_2 . Thus, embodiments of hybrid OBS network 100 enable multiple different optical cell lengths across various different carrier wavelengths. Having different optical cell lengths allows scheduler 220 to optimize the network utilization by scheduling inbound data streams 230 onto appropriately sized optical cell lengths. For instance, scheduler
20 220 may schedule a large data stream, such as data stream 230C, onto carrier wavelength λ_1 , while scheduling a small data stream, such as data stream 230D, onto carrier wavelength λ_2 . Again, the ability to schedule data slices from different data streams 230

into different optical cell lengths increases the flexibility and responsiveness of ingress node 200 to changing traffic profiles.

[0036] In a process block 340, ingress node 200 appends burst headers 280A and 280B to optical bursts 270A and 270B, respectively. A single burst header 280 is
5 transmitted per burst 270. Burst headers 280 define information pertinent to each optical burst 270, such as burst ID, burst length (i.e., number of optical cells 260 within the current optical burst), an optional header error correction (“HEC”) field, a cell loss priority (“CLP”) field, and the like. The precise header format and available fields can be tailored to the needs of the specific implementation. In some embodiments, burst
10 headers 280 are added to the first data slice of an optical burst by header pre-append block 215. In some embodiments, burst headers 280 are added to the first data slice of an optical burst by burst transmit block 225.

[0037] As can be seen from FIG. 2, stream slicer 210 must slice data blocks 235 into data slices 240 that, when appended with slice headers 245 (and in some cases
15 burst headers 280), fit into optical cells 260. Relation 1 describes the relation between the optical cell length and the data slice size,

$$OpticalCellSize \geq k \cdot (DataSliceSize + SliceHeaderSize) \quad \text{Relation 1}$$

where $k = 1, 2, 3 \dots n$. The greater the value of k , the more data slices 240 can fit into a single optical cell 260. Referring to FIG. 2, $k = 2$ for optical cells 260A of optical burst
20 270A, and $k = 1$ for optical cells 260B of optical burst 270B. It should be appreciated that k can have different values for different carrier wavelengths. Increasing k requires more expensive switching nodes 120 as they must be capable of handling larger optical

cells at a time. Thus, k may be selected based on the hardware (e.g., ingress nodes 115, egress nodes 115, and switching nodes 120) available.

5 [0038] In some embodiments, relation 1 holds true for all optical cells 260, except the first optical cell in each optical burst 270. For the first optical cell, one of the data slices must be appropriately shorter to accommodate the burst header 280. In other embodiments, an additional optical cell 260 may be appended to the front of optical bursts 270 containing only the burst header. While the latter embodiment may make the task of slicing data slices 240 less complicated for stream slicer 210, it does so at the expense of bandwidth utilization.

10 [0039] Once scheduler 220 has completely scheduled all data slices 240 into optical cells 260, then burst transmit block 225 converts the data slices 240 from the electrical realm to the optical realm and launches the optical bursts 270 into fiber link 250. Once launched, optical bursts 270 flow to one of switching nodes 120 where optical bursts 270 are processed and routed during fixed time slots, one optical cell 260
15 per carrier wavelength at a time. Eventually optical bursts 270 will reach one of egress nodes 115 for delivery to an external network and ultimately to a final destination.

[0040] Once optical bursts 270 have been launched, process 300 continues to a process block 350. In process block 350, ingress node 200 waits until the next data stream or set of data streams are buffered to repeat process 300 from process block 310.

20 [0041] FIG. 4 is a graphic illustration of process 300 for generating and transmitting fixed length optical cells 260, in accordance with an embodiment of the present invention. Stage 1 illustrates two data streams 430A and 430B. Data stream 430A is three units long while data stream 430B is two units long. Data streams 430A

and 430B are buffered within buffer 205 of ingress node 200. Stage 2 illustrates data streams 430A and 430B sliced into data slices A1-A3 and B1-B2, as described in connection with process block 325. Stage 3 illustrates data slices A1-A3 and B1-B2 having slice headers appended thereto. Finally, stage 4 illustrates data slices A1-A3 and B1-B2 scheduled into optical cells of two optical bursts 470A and 470B. Optical burst 470A includes data slices A1, B1, and A2. Optical 470B includes data slices B2, A3, and an empty cell 490. The order of the data slices A1-A3 and B1-B2 is based upon the particular scheduling algorithm implemented by scheduler 220. Although the optical cells of optical bursts 470A and 470B are illustrated as having equal lengths, it should be appreciated that the lengths of each optical burst 470A and 470B can differ.

[0042] FIG. 5 is a block diagram illustrating hardware blocks of an exemplary edge node 500, in accordance with an embodiment of the present invention. The illustrated embodiment of edge node 500 includes a processor 505, memory 510, non-volatile (“NV”) storage 515, an external network interface 520, and an internal network interface 525.

[0043] External network interface 520 provides functionality for interfacing with one or more external networks (e.g., LANs 105, WAN 107, and the like). Internal network interface 525 provides functionality for interfacing with the internal infrastructure of hybrid OBS network 100. Internal network interface 525 will typically be coupled to one or more fiber links (e.g., fiber link 250) for sending and for receiving optical bursts to/from switching nodes 120. Inbound data streams are buffered within memory 510 and processed by processor 505. In one embodiment, processor 505 is a network processor including a plurality of pipelined micro-engines capable of processing

multiple data streams in parallel. In one embodiment, each of the micro-engines corresponds to one of the functional blocks illustrated in FIG. 2 (e.g., stream slicer 210, header pre-append 215, scheduler 220, and burst transmit block 225). In some embodiments, some or all of these functional blocks are implemented in hardware. In
5 some embodiments, some or all of these functional blocks are implemented in software executed by processor 505. In one embodiment, functionality for performing operations corresponding to the process and decision blocks illustrated in FIG. 3 are stored within NV storage 515 and executed by processor 505.

[0044] The above description of illustrated embodiments of the invention,
10 including what is described in the Abstract, is not intended to be exhaustive or to limit the invention to the precise forms disclosed. While specific embodiments of, and examples for, the invention are described herein for illustrative purposes, various equivalent modifications are possible within the scope of the invention, as those skilled in the relevant art will recognize.

15 [0045] These modifications can be made to the invention in light of the above detailed description. The terms used in the following claims should not be construed to limit the invention to the specific embodiments disclosed in the specification and the claims. Rather, the scope of the invention is to be determined entirely by the following claims, which are to be construed in accordance with established doctrines of claim
20 interpretation.